

Figure 2

A recurrent neural network, consisting of a small ensemble of eight processing units. In addition to external inputs and outputs, each unit has feedback connections to other units in the network, as indicated by the branches of the output lines that loop back onto the input lines leading into each unit (reproduced by permission of the Center for the Neural Basis of Cognition from *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 2: Psychological and Biological Models*, 1986, p. 174)

the input which evokes it. This property of networks with continuous attractors may account for such peculiar effects as invariance of orientation tuning to the contrast of the stimuli (Sompolinsky and Shapley 1997), persistence of space-related activity in hippocampus in total darkness (Tsodyks 1999), and stability of gaze direction in the darkness (Seung 1996), which are difficult to explain in the more traditional, feed-forward models (see Figs. 1 and 2).

See also: Artificial Neural Networks: Neurocomputation; Cognitive Functions (Normal) and Neuropsychological Deficits, Models of; Linear Algebra for Neural Networks; Network Models of Tasks; Neural Networks and Related Statistical Latent Variable Models; Neural Networks: Biological Models and Applications; Neural Networks, Statistical Physics of; Neural Systems: Models of Behavioral Functions

Bibliography

Abeles M 1991 *Corticonics*. Cambridge University Press, New York
 Amit D J 1989 *Modeling Brain Function*. Cambridge University Press, New York

Amit D J 1994 The Hebbian paradigm reintegrated: Local reverberations as internal representations. *Behavioural Brain Science* **18**: 617–26
 Erwin E, Obermayer K, Schulten K 1995 Models of orientation and ocular dominance columns in the visual cortex: A critical comparison. *Neural Computation* **7**: 425–68
 Hebb D O 1949 *The Organization of Behavior*. Wiley, New York
 Hertz J, Krogh A, Palmer R G 1991 *Introduction to the Theory of Neural Computation*. Addison-Wesley, New York
 Hopfield J J 1982 Neural networks and physical systems with emergent collective properties. *Proceedings of the National Academy of Science USA* **79**: 2554–8
 Hopfield J J 1995 Pattern recognition computation using action potential timing for stimulus representation. *Nature* **376**: 33–6
 Kleinfeld D, Sompolinsky H 1988 Associative neural network model for the generation of temporal patterns. Theory and application to central pattern generators. *Biophysics Journal* **54**: 1039–51
 Kohonen T 1984 *Self-organization and Associative Memory*. Springer, Heidelberg, Germany
 McClelland J L, Rumelhart D E, PDP Research Group (eds.) 1986 *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. MIT Press, Cambridge, MA
 Rosenblatt F 1962 *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Spartan, Washington, DC
 Seung H S 1996 How the brain keeps the eyes still. *Proceedings of the National Academy of Science USA* **93**: 13339–44
 Sompolinsky H, Shapley R 1997 New perspectives on the mechanisms for orientation selectivity. *Current Opinion in Neurobiology* **7**: 514–22
 Tsodyks M 1999 Attractor neural network models of spatial maps in hippocampus. *Hippocampus* **9**: 481–9

M. Tsodyks

Neural Development: Mechanisms and Models

How a single cell develops into a richly structured organism, capable of processing highly complex information about the world around it, is one of the most fascinating questions of the twenty-first century. A crucial step in the development of such intelligent organisms is the formation of the nervous system. Billions of neurons must be generated, move to appropriate locations, become specialized to perform particular tasks, establish connections with the correct set of other neurons, and then sculpt their initial connectivity patterns into precisely the right configurations to subserve sophisticated information processing tasks. The last years of the twentieth century have seen a dramatic increase in our knowledge of the genes, molecules, and patterns of neural activity that control these key developmental events. Theoretical models can help organize this data, provide rigorous analyses of the key mechanisms at work, and make novel and experimentally testable predictions.

1. Overview of Neural Development

The first stage of neural development (Gilbert 2000) begins in humans three weeks after conception with the differentiation of the neural plate, a sheet of cells on the surface of the embryo. This folds in on itself to form the neural tube and the neural crest. The neural tube gives rise to the central nervous system and parts of the peripheral nervous system, while the neural crest gives rise to parts of the peripheral nervous system and some non-neural tissues. In the second stage of neural development the neural tube expands as cells proliferate and migrate to their final destinations. Cells become specified as particular types of neurons or glia by both genetic programs and factors within their environment. In stage three the initial connectivity between neurons is established. Axons elongate and are guided to their appropriate targets by a variety of molecular cues. When they reach their targets, a synapse forms between the tip of the axon and the local region of the dendrite it contacts. By the end of this stage the large-scale structure of the nervous system is established. In the fourth stage this initial crude architecture is refined by cell death, and by the gradual modification of synaptic strengths based on patterns of neural activity. These patterns of activity can be both transduced from the environment via the sense organs and generated internally via spontaneous activity. This activity molds the brain's synapses to best match the structure of the environment. It should be noted that, although these four stages have here been presented as distinct, in reality they overlap and intermingle. In addition, experimental evidence increasingly suggests that many of these processes continue to occur to some extent in the adult brain: not just the continued modification of synaptic strengths within a fixed architecture, but also the generation of new neurons, their differentiation and migration, and the formation of new connections. The majority of computational modeling work has so far been focused on understanding stage four, activity-dependent refinement of connections. There are no models addressing stage one (neural tube formation), and only very few addressing stages two (migration and differentiation), and three (axon guidance and synaptogenesis). These will be briefly discussed first, before concentrating on models of stage four.

2. Models of Pattern Formation and Regionalization

A basic problem at stage two of development is to understand how an initially spatially uniform array of cells becomes spatially nonuniform, with different neurons committed to different fates. A very influential class of mathematical models of this type of general pattern formation process are 'reaction-diffusion' systems, first proposed by Turing (1952), and subsequently developed by Meinhardt (1982). In the sim-

plest version, each cell in the array is imagined to be continually releasing two chemicals, an 'activator' and an 'inhibitor,' which diffuse freely through the array. The activator causes cells to produce both more activator and more inhibitor, while the inhibitor causes cells to produce less activator. If the inhibitor diffuses faster than the activator, the end result is the formation of a stable, non-uniform pattern consisting of small regions with high concentrations of activator surrounded by regions with high concentrations of inhibitor. Although this is mathematically elegant and has the potential to explain many aspects of biological pattern formation (Meinhardt 1982 Murray 1993), there are actually very few biological systems where reaction-diffusion has been definitively established as the pattern formation mechanism employed. In fact, experimental evidence increasingly suggests that a rather different process may be more important for regional specification of the nervous system. Homologues of many of the genes that control nervous system regionalization in mammals were first identified in the formation of the segmented bodies of fruit flies. Recent molecular data suggests that, in fruit flies, an initial anterior-posterior molecular gradient is converted into differential levels of expression of particular genes via different thresholds for their activation. This leads to differential levels of expression of further genes, and so on until a complex mosaic of regions is formed. This was first discussed theoretically as a mechanism for regionalization by Wolpert (1969), and it seems highly likely that these types of mechanisms are at work in the regionalization of the mammalian nervous system (Rubenstein and Beachy 1998) rather than reaction-diffusion. Computational modeling has only recently begun to address how complex networks of interacting proteins can lead to the formation of spatial patterns (e.g., Reinitz and Sharp 1995).

3. Models of Axon Guidance

A basic problem at stage three of development is to understand how an extending axon navigates over long distances through a complex molecular environment to find an appropriate target. Among the cues involved are molecules that promote and inhibit axonal growth, and molecules that pull axons in particular directions when they are present in a concentration gradient (Tessier-Lavigne and Goodman 1996). Axons detect the distribution of molecules in their environment via a dynamic sensing device at their tip called the growth cone (see Fig. 1). This draws the axon up or down gradients of molecules such as netrin-1 by sensing a concentration difference across its spatial extent. Crucial to this process are filopodia (see Fig.1), long and dynamic protrusions from the growth cone. Modeling work has approached the growth cone in several different ways, attempting to understand the minimum gradient steepness that

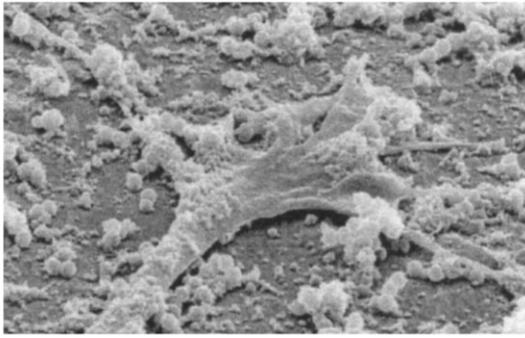


Figure 1

Electron micrograph of a growth cone at the end of an axon (here the growth cone is growing over a surface irregularly covered with coated beads). The long finger-like protrusions are filopodia. Growth cones are typically about 0.01 mm across. (Source: Rosentreter et al. 1998, *Journal of Neurobiology* 37 541–62, ©John Wiley and Sons)

growth cones can detect (Goodhill and Urbach 1999), movement of filopodia (Buettner 1995), how filopodia could be generated as the result of a reaction–diffusion mechanism inside the growth cone (Meinhardt 1999), and how molecules released by growth cones themselves could help guide them to their targets (Hentschel and Van Ooyen 1999). Besides modeling the growth cone, computational researchers have also considered the gradient shapes that could be present *in vivo* and how these shapes constrain the maximum distance over which axons could be guided (Goodhill 1998). A very popular system for studying axon guidance by gradients is the growth of retinal axons into the optic tectum to form a topographic map. Sperry (1963) suggested that this process could be implemented by a matching process between molecular gradients in the retina and the tectum. Following Sperry many experiments were performed to test this hypothesis, and a number of theoretical models were proposed based on concepts of both gradient matching and competition between retinal axons for space in the tectum (e.g., Prestige and Willshaw 1975). More recently gradients of Eph receptors have been identified in the retina, and gradients of their ligands the ephrins in the tectum, which could be a molecular substrate for Sperry's hypothesis (Flanagan and Vanderhaeghen 1998). So far, however, theoretical models have not kept pace with these developments (Goodhill and Richards 1999).

4. Models of Activity-dependent Refinement of Connections

A basic problem at stage four of development is to understand how the initial pattern of synaptic

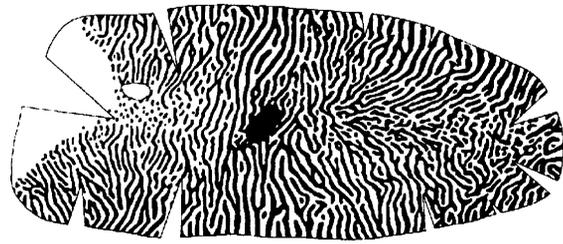


Figure 2

The pattern of ocular dominance columns from the primary visual cortex of a macaque monkey. White represents regions of cortex dominated by input from one eye, black the other eye. Each individual (including each human) apparently has a different ocular dominance pattern that is unique as a fingerprint. The width of individual columns is 0.5–1 mm. (Source: Levay et al. 1985 *Journal of Neuroscience* 5 486–501, © Society for Neuroscience. Reprinted by permission of Wiley–Liss, Inc., a Subsidiary of John Wiley and Sons, Inc.)

strengths established by stage three of development is refined and sculpted by neural activity. It is to this stage that theoretical modeling has so far been most extensively applied, perhaps because of the link between experimentally determined mechanisms of synaptic plasticity and learning algorithms for artificial neural networks (Hertz et al. 1991). The key organizing principle is the Hebb rule (Hebb 1949), which can be summarized as ‘cells that fire together wire together.’ In other words, the strength of a synapse tends to increase when the firing of the presynaptic neuron is correlated with the firing of the postsynaptic neuron. Hebb's qualitative statement can be made mathematically precise in several different ways.

Perhaps the paradigm example of the application of Hebbian learning rules to understanding neural development is modeling of the development of receptive fields and maps in the mammalian visual system. Visual information passes initially from the retina, via the lateral geniculate nucleus (LGN), to the primary visual cortex (V1). Receptive fields of neurons in V1 can be crudely characterized as preferring (being most strongly excited by) input from one small region of visual space, more from one eye than the other eye (‘ocular dominance’), and in the form of an edge or bar of light of a particular orientation (Hubel 1988). Nearby neurons in V1 usually prefer similar positions in space, similar orientations, and have a similar eye preference, but there are also occasional discontinuities where these preferences change abruptly. This leads receptive field properties to be laid out across the cortex in the form of maps with very particular spatial properties (a map of ocular dominance is shown in Fig. 2). These properties of V1 have been appealing to theoretical modelers because they seem to progress to this final, highly organized state from a far less

organized initial state, and this process of 'self-organization' depends on both the presence and correlational structure of neural activity in the retina in a way broadly consistent with the Hebb rule. The first mathematical model of receptive field development in V1 was put forward by von der Malsburg (1973), and since then a large number of variations on the Hebbian theme have been proposed (reviewed in Swindale 1996). A key element of all of the models addressing map formation is the assumption of some kind of lateral interaction between nearby cortical neurons. This is usually short-range excitation balanced against longer-range inhibition, which gives these models a mathematically somewhat analogous form to the reaction-diffusion models that have been applied to spatial pattern formation at much earlier stages of development.

A key feature of these visual maps is their plasticity in response to the correlational structure of input from the environment. The paradigm example is the change in the ocular dominance map when one eye is deprived of visual input. In this case the regions of the cortex dominated by the open eye expand at the expense of the regions dominated by the closed eye (Hubel 1988 Katz and Shatz 1996). These and other experiments indicate that axons from the two eyes compete for space in the primary visual cortex. This can be modeled theoretically in a similar way to the competition for space between individual retinal axons in topographic map formation in the optic tectum (Goodhill 1993). In other words, there are active processes by which the developing nervous system adapts its structure to make the best use of the available environmental input. This also shows up in, for instance, increases in the size of the cortical representation of body surfaces that are highly stimulated (Buonomano and Merzenich 1998). Interestingly, much of this kind of plasticity in the nervous system is limited to certain temporal windows in development called critical periods. For instance, the critical period for the plasticity of ocular dominance columns in kittens is between three and six weeks after birth: changing the visual input before or after this time has relatively little effect. Critical periods also exist for the development of other parts of the nervous system, but what controls the timing of critical periods is unknown. Possibilities include a simple genetically-controlled gating mechanism that allows plasticity only for a limited period, or a less rigid mechanism whereby the dynamics of the learning algorithm naturally lead to a state that is relatively stable, as in some theoretical models (e.g. Linsker 1986).

However, despite the presence of critical periods, the structure of the nervous system continues to evolve for a very long time after the initial processes of development are complete. For instance, a recent study has shown continued growth in the brains of children up to at least 15 years old (Thompson et al. 2000). Although it used to be thought that the neurons

born initially during development are the total complement available, recent research indicates that new neurons may be being born in specialized regions and integrated into the nervous system throughout life (Gage 2000). Theoretical modeling has barely begun to address such phenomena. Neural development will clearly continue to be an expanding frontier of scientific enquiry for some time to come.

See also: Artificial Neural Networks: Neurocomputation; Brain Development, Ontogenetic Neurobiology of; Cognitive Development in Infancy: Neural Mechanisms; Computational Neuroscience; Connectionist Models of Development; Hebb, Donald Olding (1904–85); Neural Networks: Biological Models and Applications; Neural Plasticity; Neural Plasticity in Auditory Cortex; Neural Plasticity in Visual Cortex; Neural Systems: Models of Behavioral Functions; Prenatal and Infant Development: Overview

Bibliography

- Buettner H M 1995 Computer simulation of nerve growth cone filopodial dynamics for visualization and analysis. *Cell Motility and the Cytoskeleton* **32**: 187–204
- Buonomano D V, Merzenich M M 1998 Cortical plasticity: From synapses to maps. *Annual Review of Neuroscience* **21**: 149–86
- Flanagan J G, Vanderhaeghen P 1998 The ephrins and eph receptors in neural development. *Annual Review of Neuroscience* **21**: 309–45
- Gage F H 2000 Mammalian neural stem cells. *Science* **287**: 1433–8
- Gilbert S F 2000 *Developmental Biology*. Sinauer, Sunderland, MA
- Goodhill G J 2000 Topography and ocular dominance: A model exploring positive correlations. *Biological Cybernetics* **69**: 109–18
- Goodhill G J 1998 Mathematical guidance for axons. *Trends in Neurosciences* **21**: 226–31
- Goodhill G J, Richards L J 1999 Retinotectal maps: Molecules, models, and misplaced data. *Trends in Neurosciences* **22**: 529–34
- Goodhill G J, Urbach J S 1999 Theoretical analysis of gradient detection by growth cones. *Journal of Neurobiology* **41**: 230–41
- Hebb D O 1949 *The Organization of Behaviour*. Wiley, New York
- Hentschel H G E, Van Ooyen A 1999 Models of axon guidance and bundling during development. *Proceedings of the Royal Society of London B* **266**: 2231–8
- Hertz J, Krogh A, Palmer R G 1991 *Introduction to the Theory of Neural Computation. Lecture Notes in the Santa Fe Institute Studies in the Sciences of Complexity*. Addison-Wesley, Redwood City, CA
- Hubel D H 1988 *Eye, Brain and Vision*. Freeman, New York
- Katz L C, Shatz C J 1996 Synaptic activity and the construction of cortical circuits. *Science* **274**: 1133–38
- Levay, Connolly M, Houde J, Van Essen D C 1985 *Journal of Neuroscience* **5**: 486–501

- Linsker R 1986 From basic network principles to neural architecture. *Proceedings of the National Academy of Sciences USA* **83**: 7508–12, 8390–4, 8779–83
- Malsburg C von der 1973 Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik* **14**: 85–100
- Meinhardt H 1982 *Models of Biological Pattern Formation*. Academic Press, London
- Meinhardt H 1999 Orientation of chemotactic cells and growth cones: models and mechanisms. *Journal of Cell Science* **112**: 2867–74
- Murray J D 1993 *Mathematical Biology*, 2nd edn. Springer, Berlin
- Prestige M C, Willshaw D J 1975 On a role for competition in the formation of patterned neural connexions. *Proceedings of the Royal Society of London B* **190**: 77–98
- Reinitz J, Sharp D H 1995 Mechanism of *eve* stripe formation. *Mechanisms of Development* **49**: 133–158
- Rosentreter S M, Davenport R W, Löschinger J, Huf J, Jung J, Bonhoeffer F 1998 *Journal of Neurobiology* **37**: 541–62
- Rubenstein J L, Beachy P A 1998 Patterning of the embryonic forebrain. *Current Opinions in Neurobiology* **8**: 18–26
- Sperry R W 1963 Chemoaffinity in the orderly growth of nerve fiber patterns and connections. *Proceedings of the National Academy of Sciences USA* **50**: 703–10
- Swindale N V 1996 The development of topography in the visual cortex: A review of models. *Network* **7**: 161–247
- Tessier-Lavigne M, Goodman C S 1996 The molecular biology of axon guidance. *Science* **274**: 1123–33
- Thompson P M, Giedd J N, Woods R P, MacDonald D, Evans A C, Toga A W 2000 Growth patterns in the developing brain detected by using continuum mechanical tensor maps. *Nature* **404**: 190–3
- Turing A M 1952 The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London B* **237**: 37–72
- Wolpert L 1969 Positional information and the spatial pattern of cellular differentiation. *Journal of Theoretical Biology* **25**: 1–47

G. J. Goodhill

Neural Networks and Related Statistical Latent Variable Models

1. Introduction

This article presents statistical features associated with artificial neural networks, as well as two other related models, namely mixtures-of-experts and hidden Markov models.

Historically, the study of artificial neural networks was motivated by attempts to understand information processing in biological nervous systems (McCulloch and Pitts 1943). Recognizing that the mammalian brain consists of very complex webs of interconnected neurons, artificial neural networks are built from densely interconnected simple processing units, where each unit receives a number of inputs and produces a single numerical output. However, while artificial neural networks have been motivated by biological systems, there are many aspects of biological systems not adequately modeled by artificial neural networks.

Indeed, many aspects of artificial neural networks are inconsistent with biological systems. Though the brain metaphor provides a useful source of inspiration, the viewpoint adopted here is that artificial neural networks are a general class of parameterized statistical models consisting of interconnected processing units.

Artificial neural networks provide a general and practical approach for approximating real-valued (regression problems) and discrete-valued (classification problems) mappings between covariate and response variables. These models, combined with simple methods of estimation (i.e., learning algorithms), have proven highly successful in such diverse areas as handwriting recognition (LeCun et al. 1989), spoken word recognition (Lang et al. 1990), face recognition (Cottrell 1990), text-to-speech translation (Sejnowski and Rosenberg 1986), and autonomous vehicle navigation (Pomerleau 1993).

An artificial neural network consists of multiple sets of units: one set of units corresponds to the covariate variables, a second set of units corresponds to the response variables, and a third set of units, referred to as hidden units, corresponds to latent variables. The hidden units (or, more precisely, their associated hidden or latent variables) mediate the nonlinear mapping from covariate to response variables. Based on observed pairings of covariate and response variables, the back-propagation algorithm is commonly used to estimate the parameters, called weights, of the hidden units.

Some mixture models have statistical structures that resemble those of artificial neural networks. Mixture models are multiple-component models in which each observable data item is generated by one, and only one, component of the model. Like artificial neural networks, mixture models also contain latent or hidden variables. The hidden variables of a mixture model indicate the component that generated each data item.

Mixtures-of-experts (ME) models combine simple conditional probability distributions, such as unimodal distributions, in order to form a complex (e.g., multimodal) conditional distribution of the response variables given the covariate variables. They combine properties of generalized linear models (McCullagh and Nelder 1989) with those of mixture models. Like generalized linear models, they are used to model the relationship between covariate and response variables. Typical applications include nonlinear regression and binary or multiway classification. However, unlike standard generalized linear models they assume that the conditional distribution of the responses is a finite mixture distribution. Mixtures-of-experts are 'piecewise estimators' in the sense that different mixture components summarize the relationship between covariate and response variables for different subsets of the observable data items. The subsets do not, however, have hard boundaries; a data item might simultaneously be a member of multiple subsets. Because

10526